Diff-SABR: Diffusion-Based Symmetry-Aware Reconstruction of Alveolar Bone Defects from CBCT

Ho Yoon Choi[1*], Ji Yong Han[2*], Sang-Heon Lim[2], Sujeong Kim[2], Hyunbin Yun[2], Mobin Ahmadi[2], Sungho Oh[2], Taeyeon Lee[2], Jeong Joon Han[3], and Won-Jin Yi[1,2,4]

[1] Applied Bioengineering, Graduate School of Convergence Science and Technology, Seoul National University, Republic of Korea
[2] Interdisciplinary Program in Bioengineering, Graduate School of Engineering, Seoul National University, Republic of Korea
[3] Oral and Maxillofacial Surgery and Dental Research Institute, School of Dentistry, Seoul National University, Republic of Korea
[4] Oral and Maxillofacial Radiology and Dental Research Institute, School of Dentistry, Seoul National University, Republic of Korea
wjyi@snu.ac.kr

Abstract. Precise 3D reconstruction of alveolar bone defects using Cone-Beam Computed Tomography (CBCT) is essential for predictable dental implant surgery and bone graft planning. However, automated restoration remains challenging due to the inherent anatomical complexity and the scarcity of paired clinical datasets. In this paper, we propose Diff-SABR, a novel diffusion-based symmetry-aware framework to reconstruct alveolar bone morphology. Our approach employs a hybrid learning strategy that integrates a synthetic dataset generated through hyperbolic-paraboloid simulation with a clinical dataset for weakly supervised training. To handle anatomical variations, we incorporate AlignNet, a spatial alignment module that registers healthy mirrored anatomy to the defect site. Experimental results demonstrate that our framework achieves high anatomical fidelity, with a Dice Similarity Coefficient (DSC) of 0.82 and a 95th percentile Hausdorff Distance (HD95) of 1.67 mm on real-world clinical scans. Quantitative and qualitative ablation studies further confirm that the synergy between hybrid supervision and spatial alignment is crucial to restore clinically viable ridge geometries. By providing objective and automated 3D bone restoration, Diff-SABR offers a robust tool to improve surgical stability and optimize bone graft material selection in clinical practice.

Keywords: Diff-SABR · Alveolar Bone Reconstruction · Latent Diffusion Model · Symmetry-aware Learning · Weakly Supervised Learning

* These authors contributed equally.

1 Introduction

Precise 3D reconstruction of alveolar bone defects using Cone-Beam Computed Tomography (CBCT) is a prerequisite for successful dental implant surgery [3, 14, 18]. Post-extraction bone resorption often requires bone grafting, making quantitative assessment of defects essential for determining the graft volume and ensuring long-term surgical stability [17]. However, clinical evaluation still relies largely on subjective judgment due to the lack of objective and automated restoration tools [5, 8, 19]. Early automation efforts, such as intensity-based thresholding [14] and template-driven completion [11, 2], are often hampered by inter-operator variability and high sensitivity to alignment. Furthermore, the low contrast of CBCT and the structural diversity of defects restrict the robustness of these intensity-driven priors in complex clinical scenarios [13, 15].

Motivated by these limitations, recent studies have explored deep generative models for anatomically consistent reconstruction [6, 12, 20]. Diffusion probabilistic models, in particular, have emerged as powerful conditional generative frameworks capable of synthesizing high-fidelity structures [16, 10]. Diffusion models progressively transform noise into anatomically consistent volumes conditioned on auxiliary priors, making them well-suited for CBCT-based restoration. Despite this potential, automated restoration remains challenging due to anatomical complexity and the scarcity of paired clinical datasets for supervised learning. While simulation-based training is widely adopted to address this limitation, existing approaches relying on naive masking or random cropping fail to capture the irregular and asymmetric nature of bone resorption. Stochastic anomaly simulation for volumetric inpainting [9] further lacks the structural realism required to model complex multi-site bone loss.

To address these challenges, we propose Diff-SABR (Diffusion-based Symmetry-aware Bone Reconstruction). Diff-SABR introduces a hyperbolic-paraboloid-driven defect simulation strategy that enables the synthesis of geometrically principled resorption patterns across the maxillofacial complex. Taking advantage of the bilateral symmetry of the jawbone, our framework enforces latent alignment between defect regions and contralateral healthy structures to promote global morphological consistency. Additionally, we incorporate a weakly supervised learning objective to bridge the gap between the simulated and clinical domains without requiring dense voxel-wise clinical annotations. Our contributions are summarized as follows:

- We introduce a hyperbolic-paraboloid-driven simulation strategy that mathematically replicates asymmetric bone resorption patterns, providing a principled foundation for training without paired clinical datasets.
- We propose Diff-SABR, a symmetry-aware conditional diffusion framework integrating manifold-constrained latent alignment with a weakly supervised objective to ensure anatomical plausibility across domains.
- We demonstrate the framework's superior performance in volumetric quantification for dental implantology through evaluations on both synthetic and clinical datasets.

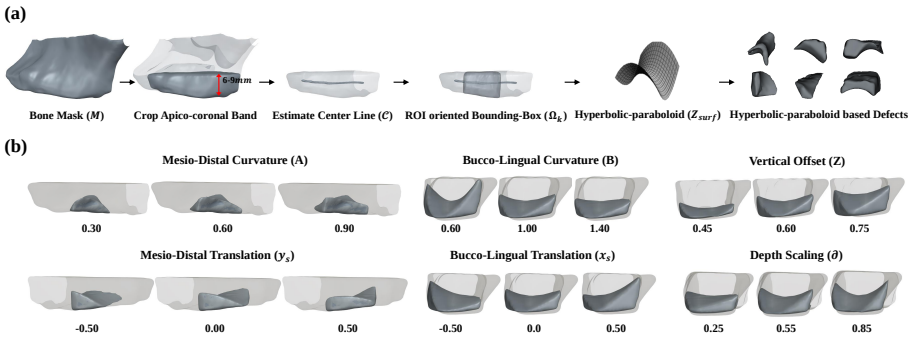

Fig. 1. Hyperbolic-paraboloid defect simulation. (a) Synthetic data generation pipeline and (b) defect variations generated by modulating geometric parameters θ.

2 Method

2.1 Hyperbolic-Paraboloid-Driven Defect Simulation

We generate clinically plausible paired data by simulating resorption patterns in intact jawbone CBCT masks (Fig. 1(a)). Given a segmented bone mask M, an apico-coronal band of 6-9 mm is cropped to approximate the clinically relevant ridge region for augmentation planning [3, 18]. To realistically position defects along the dental arch, we estimate a ridge centerline $\mathcal{C}$ using Principal Component Analysis (PCA) and randomly sample a center point $\mathbf{c}_k \in \mathcal{C}$. Around $\mathbf{c}_k$, we define a local coordinate system aligned with the arch tangent (mesio-distal) and its orthogonal bucco-lingual direction, constructing an oriented bounding-box Ω_k. Within this ROI, we synthesize a saddle-shaped defect surface using a hyperbolic paraboloid:

$$Z_{\text{surf}}(x_n, y_n; \theta) = Z + \partial \left(\frac{(y_n - y_s)^2}{A^2} - \frac{(x_n - x_s)^2}{B^2} \right), \tag{1}$$

where $\theta = \{Z, \partial, A, B, x_s, y_s\}$ controls the offset, depth, curvature, and asymmetric shifts in the local plane (x_n, y_n) (Fig. 1(b)). Finally, we remove bone voxels inside Ω_k based on Z_{surf} to generate the defective mask M_{def}, defining the target defect region as $D = M \setminus M_{\text{def}}$. By repeating this procedure with stochastic variations of θ, we produce a diverse and anatomically consistent set of defect patterns for diffusion-based training.

2.2 Diff-SABR: Symmetry-aware Bone Reconstruction Diffusion Framework

Network Overview. The Diff-SABR framework reconstructs alveolar bone defects by integrating structural priors through a two-stage learning process [16].

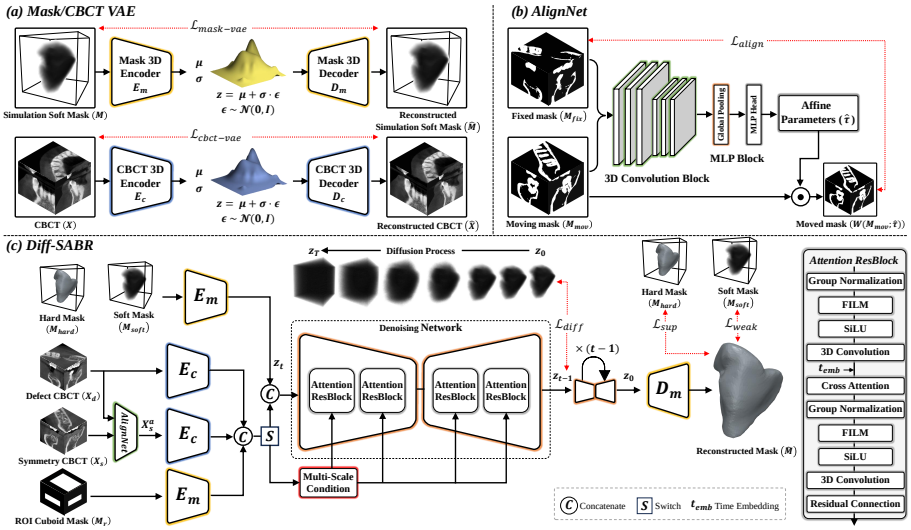

Fig. 2. Overview of Diff-SABR, including pre-trained VAEs (a), AlignNet (b), and a symmetry-aware latent diffusion model (c).

The initial stage focuses on establishing a robust latent space via pre-trained VAEs and achieving spatial consistency through an alignment module. This is followed by a fine-tuning stage where a symmetry-aware latent diffusion model performs the actual bone restoration within the learned manifold. By combining synthetic data for full supervision and clinical data for weak supervision, the network effectively bridges the domain gap while preserving high anatomical fidelity. The following sections detail each component and the hybrid optimization strategy.

Pre-training Stage. The framework begins with the pre-training of three modules to establish a robust latent representation and spatial alignment. Two independent 3D VAEs are optimized using reconstruction and KL divergence terms to learn compact manifolds for CBCT intensities and binary or soft masks, respectively. The resulting CBCT encoder E_c and Mask encoder-decoder pair E_m, D_m facilitate stable diffusion by projecting high-dimensional volumes into a low-dimensional latent space. Simultaneously, AlignNet is trained to regress an affine transform $\hat{\tau}$ that registers the mirrored contralateral anatomy to the defect site, thereby reducing pose-induced misalignments. For AlignNet, the fixed and moving masks denote thresholded bone-only maps extracted from the defect-side CBCT and the mirrored contralateral CBCT, respectively, and are used only to estimate the affine symmetry alignment. The aligned symmetry CBCT is then obtained as $X_s^a = W(X_s; \hat{\tau})$, ensuring reliable symmetry-based conditioning for the subsequent reconstruction.

Symmetry-aware Latent Diffusion. During the fine-tuning stage, we optimize a conditional diffusion model that operates within the pre-trained latent

manifold. Since synthetic and clinical data provide different types of supervision, we define the target mask M accordingly. For synthetic data, the hard mask M_{hard} is the binary defect region generated by the hyperbolic-paraboloid simulation, corresponding to the removed bone region between the intact and simulated defective masks. For clinical data, where voxel-wise ground truth is unavailable, the discrepancy between the defect-side anatomy and the aligned mirrored contralateral reference within the clinical defect ROI is converted into a continuous soft representation using a signed distance function (SDF), and used as the clinical pseudo label M_{soft} for weak supervision. During training, M_{hard} provides full supervision, whereas M_{soft} provides weak anatomical guidance. In contrast to these supervision masks, the ROI cuboid mask M_r serves only as a coarse localization prior for the candidate defect region.

This target is encoded into the latent space as $z_0 = E_m(M)$ before applying the forward diffusion process $z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1-\bar{\alpha}_t}\epsilon$, where $\epsilon \sim \mathcal{N}(0, I)$. For conditional guidance, we construct a joint feature $\mathcal{C}_0$ by concatenating the embeddings of the defect CBCT, the aligned symmetry CBCT, and the ROI cuboid mask, formulated as $\mathcal{C}_0 = \text{Concat}(E_c(X_d), E_c(X_s^a), E_m(M_r))$.

To ensure consistent guidance across different spatial resolutions, the condition feature $\mathcal{C}_0$ is transformed into a multi-scale pyramid through iterative average pooling. At the input layer, the noisy latent z_t and the full-resolution condition $\mathcal{C}_0$ are concatenated and processed through an initial 3D convolutional block Φ_{in}. These multi-scale features are subsequently injected into the encoder and decoder blocks of the denoising network using FiLM modulation and cross-attention mechanisms [15, 10]. This integrated approach allows the network to effectively fuse auxiliary priors with latent activations for high-fidelity bone restoration. The model is optimized through a total weighted loss that balances diffusion accuracy with anatomical supervision. The primary diffusion loss $\mathcal{L}_{\text{diff}} = \mathbb{E}_{t,\epsilon}[\|\epsilon - \epsilon_\phi(z_t, t, \mathcal{C}_0)\|_2^2]$ ensures the fidelity of the generative process. For synthetic data, we enforce anatomical accuracy using $\mathcal{L}_{\text{sup}} = \mathcal{L}_{\text{BCE}} + \mathcal{L}_{\text{Dice}}$ between the predicted mask $\hat{M}$ and the simulation label M_{hard}. For clinical data where pixel-wise ground truth is unavailable, the weak supervision loss is defined as

$$\mathcal{L}_{\text{weak}} = \begin{cases} 0.5(\hat{M} - M_{soft})^2/\beta, & \text{if } |\hat{M} - M_{soft}| < \beta \\ |\hat{M} - M_{soft}| - 0.5\beta, & \text{otherwise} \end{cases} \tag{2}$$

where β denotes the threshold for the Huber-like penalty. The final objective is

$$\mathcal{L}_{\text{total}} = \lambda_{\text{diff}}\mathcal{L}_{\text{diff}} + \lambda_{\text{sup}}\mathcal{L}_{\text{sup}} + \lambda_{\text{weak}}\mathcal{L}_{\text{weak}} \tag{3}$$

where the weights are empirically set to $\lambda_{\text{diff}} = 1$, $\lambda_{\text{sup}} = 1$, and $\lambda_{\text{weak}} = 0.5$.

3 Experiments

3.1 Datasets and Implementation Details

Dataset. The study utilized a CBCT dataset consisting of 33 normal subjects with complete dentition for defect simulation and 43 clinical cases featuring alveolar bone defects from patients requiring bone grafting and implant surgery.

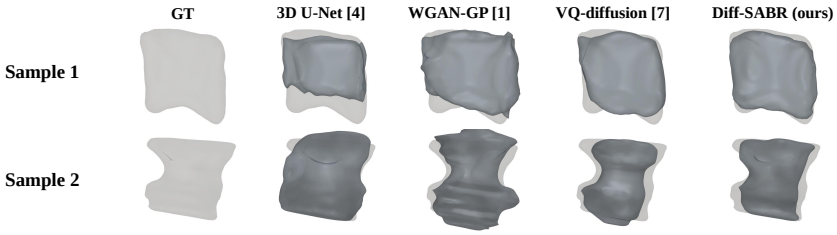

Fig. 3. Qualitative comparison of ground truth surfaces (gray) and reconstructed defect surfaces (dark gray).

All scans were acquired using a CS 9600 scanner (Carestream Health, Inc., Rochester, US) at 90 kVp and 8 mA with an isotropic voxel size of $0.15 \times 0.15 \times 0.15$ mm^3 and dimensions of 1067×1067 pixels. For the simulated dataset $\mathcal{D}_s$, 230 normal ROI volumes of size $128 \times 128 \times 128$ voxels were extracted around the alveolar ridge. Three defect patterns were generated per ROI using hyperbolic-paraboloid simulation with stochastic parameters θ, resulting in 690 samples divided into 641 for training and 49 for evaluation. The clinical dataset $\mathcal{D}_c$ was similarly cropped into ROIs centered on defects where 33 cases were assigned to weakly supervised training and 10 cases were reserved for testing.

Training setup. All models were implemented in PyTorch and trained on an NVIDIA RTX A6000 GPU (48 GB memory) under identical computing environments using a batch size of 2 to ensure a fair comparison. AlignNet and 3D VAE were optimized using the Adam optimizer with a learning rate of 1×10^{-4} and trained for 50 and 300 epochs. Diff-SABR was trained using the AdamW optimizer with a learning rate of 5×10^{-5}. A forward diffusion process with $T = 1000$ timesteps was employed, using a linear beta schedule ranging from $1e-4$ to 0.002.

Evaluation Metrics. To evaluate reconstruction accuracy, we quantified overlap-based performance using Dice Similarity Coefficient (DSC), Precision, and Recall, alongside surface-based distance metrics, including the 95th percentile Hausdorff Distance (HD95) and Mean Surface Distance(MSD).

3.2 Experimental Results

Quantitative comparison with other methods. We compare Diff-SABR with a supervised 3D U-Net [4], a WGAN-GP-based generative inpainting model [1], and a discrete-latent VQ-diffusion baseline [7]. As reported in Table 1, Diff-SABR achieves competitive performance on the simulated dataset $\mathcal{D}_s$ while producing the lowest surface distance errors (HD95 and MSD), indicating improved geometric fidelity. On the clinical dataset $\mathcal{D}_c$, Diff-SABR achieves the best overall performance, with the highest DSC, precision, and recall together with reduced surface errors. The consistent improvement on $\mathcal{D}_c$ indicates that the proposed symmetry-aware conditioning and weakly supervised adaptation effectively mit-

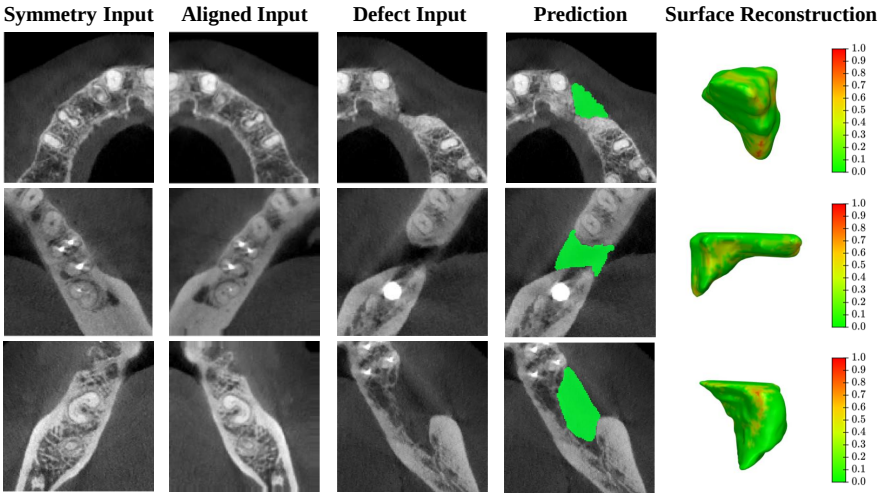

Fig. 4. Clinical CBCT reconstruction examples in axial slices. Each row corresponds to a different case. Color maps indicate surface distance errors (mm).

Table 1. Quantitative comparison on synthetic ($\mathcal{D}_s$) and clinical ($\mathcal{D}_c$) datasets.

Method	$\mathcal{D}_s$					$\mathcal{D}_c$				
	DSC (↑)	Precision (↑)	Recall (↑)	HD95 (↓)	MSD (↓)	DSC (↑)	Precision (↑)	Recall (↑)	HD95 (↓)	MSD (↓)
3D U-Net[4]	$\mathbf{0.95 \pm 0.05}$	$\mathbf{0.92 \pm 0.08}$	$\mathbf{0.98 \pm 0.01}$	2.27 ± 2.13	0.19 ± 0.19	0.65 ± 0.06	0.69 ± 0.12	0.64 ± 0.11	2.58 ± 0.94	0.72 ± 0.18
WGAN-GP[1]	0.81 ± 0.09	0.85 ± 0.14	0.80 ± 0.06	3.99 ± 2.75	0.51 ± 0.31	0.74 ± 0.04	0.77 ± 0.11	0.73 ± 0.09	1.99 ± 0.83	0.51 ± 0.13
VQ-diffusion[7]	0.87 ± 0.07	0.87 ± 0.10	0.88 ± 0.05	1.87 ± 2.03	0.27 ± 0.18	0.76 ± 0.04	0.76 ± 0.08	0.77 ± 0.08	2.22 ± 0.91	0.48 ± 0.13
Diff-SABR	0.91 ± 0.03	0.90 ± 0.05	0.91 ± 0.05	$\mathbf{0.96 \pm 0.68}$	$\mathbf{0.16 \pm 0.11}$	$\mathbf{0.82 \pm 0.04}$	$\mathbf{0.83 \pm 0.08}$	$\mathbf{0.82 \pm 0.06}$	$\mathbf{1.67 \pm 0.76}$	$\mathbf{0.36 \pm 0.08}$

igate the synthetic-to-clinical domain gap, which remains challenging for purely supervised or purely generative baselines.

Qualitative Analysis. Qualitative comparisons on clinical CBCT scans demonstrate that 3D U-Net often produces over-smoothed reconstructions and fails to recover plausible ridge contours when defect geometries deviate from simulation patterns (Fig. 3). WGAN-GP and VQ-diffusion improve continuity but often generate irregular boundaries, particularly in thin cortical regions. In contrast, Diff-SABR reconstructs coherent ridge morphology with smooth transitions at defect boundaries, yielding anatomically consistent masks that align closely with reference shapes. The integration of the symmetry prior is further detailed in Fig. 4. Due to anatomical asymmetry, the symmetry CBCT is initially misaligned with the defect site, and AlignNet transforms it into a registered structural reference. The aligned volume serves as a conditioning prior for diffusion, producing reconstructions with consistent boundaries and anatomically plausible curvature.

Effectiveness of Hybrid Supervision and Alignment Modules. Table 2 summarizes the ablation study. Training with $\mathcal{L}_{\text{sup}}$ alone achieves strong performance on $\mathcal{D}_s$ with DSC 0.91 ± 0.05 but generalizes poorly to $\mathcal{D}_c$ with DSC 0.04 ± 0.06 with large surface errors, confirming a severe domain gap. Training with $\mathcal{L}_{\text{weak}}$ alone improves $\mathcal{D}_c$ to DSC 0.56 ± 0.10 but performs poorly on $\mathcal{D}_s$

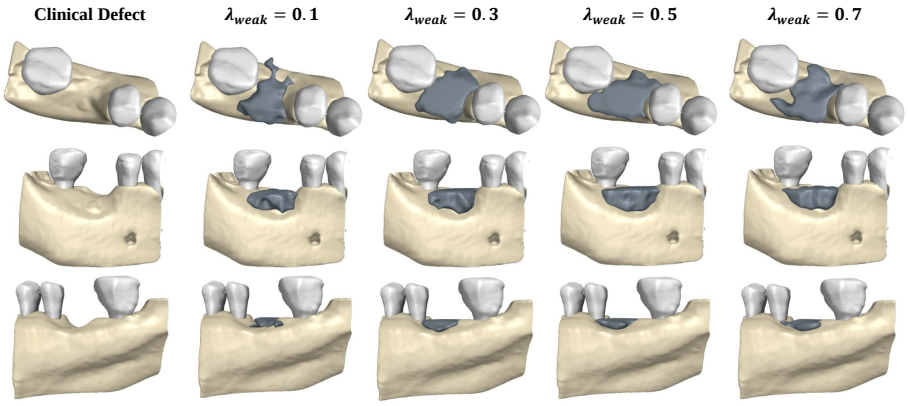

Fig. 5. Qualitative comparison of λ_{weak} from 0.1 to 0.7 in axial, buccal, and lingual views.

Table 2. Ablation study on synthetic ($\mathcal{D}_s$) and clinical ($\mathcal{D}_c$) datasets.

Components			$\mathcal{D}_s$					$\mathcal{D}_c$				
$\mathcal{L}_{sup}$	$\mathcal{L}_{weak}$	AlignNet	DSC (↑)	Precision (↑)	Recall (↑)	HD95 (↓)	MSD (↓)	DSC (↑)	Precision (↑)	Recall (↑)	HD95 (↓)	MSD (↓)
✓			0.91 ± 0.05	0.91 ± 0.05	0.91 ± 0.05	0.87 ± 0.89	0.15 ± 0.06	0.04 ± 0.06	0.36 ± 0.37	0.02 ± 0.03	8.30 ± 2.82	3.36 ± 1.37
	✓		0.25 ± 0.12	0.20 ± 0.13	0.42 ± 0.12	7.36 ± 1.42	1.84 ± 0.44	0.56 ± 0.10	0.59 ± 0.18	0.59 ± 0.14	3.72 ± 1.16	0.98 ± 0.21
✓		✓	$\mathbf{0.92 \pm 0.03}$	$\mathbf{0.92 \pm 0.04}$	0.91 ± 0.06	$\mathbf{0.76 \pm 0.69}$	$\mathbf{0.14 \pm 0.05}$	0.13 ± 0.21	0.58 ± 0.38	0.10 ± 0.18	7.17 ± 2.88	2.37 ± 1.31
	✓	✓	0.24 ± 0.12	0.19 ± 0.12	0.35 ± 0.11	7.32 ± 1.47	1.88 ± 0.47	0.57 ± 0.10	0.57 ± 0.14	0.60 ± 0.14	3.54 ± 0.84	0.92 ± 0.22
✓	✓		0.90 ± 0.06	0.86 ± 0.09	$\mathbf{0.94 \pm 0.03}$	1.00 ± 0.97	0.17 ± 0.09	0.72 ± 0.06	0.76 ± 0.11	0.70 ± 0.09	2.41 ± 1.02	0.58 ± 0.21
✓	✓	✓	0.91 ± 0.03	0.90 ± 0.05	0.91 ± 0.05	0.96 ± 0.68	0.16 ± 0.11	$\mathbf{0.82 \pm 0.04}$	$\mathbf{0.83 \pm 0.08}$	$\mathbf{0.82 \pm 0.06}$	$\mathbf{1.67 \pm 0.76}$	$\mathbf{0.36 \pm 0.08}$

with DSC 0.25 ± 0.12, indicating that weak clinical cues alone are insufficient to learn robust anatomical completion. Combining $\mathcal{L}_{\text{sup}}$ and $\mathcal{L}_{\text{weak}}$ improves clinical reconstruction to DSC 0.72 ± 0.06 and HD95 2.41 ± 1.02, demonstrating that simulation supervision and weak clinical adaptation are complementary. Adding AlignNet further boosts clinical performance to DSC 0.82 ± 0.04 and reduces HD95 to 1.67 ± 0.76 and MSD to 0.36 ± 0.08, demonstrating the benefit of explicit symmetry alignment.

Sensitivity analysis on λ_{weak}. We evaluated the effect of the weak supervision weight by varying $\lambda_{\text{weak}} \in \{0.1, 0.3, 0.5, 0.7\}$ as presented in Table 3. Clinical performance peaked at $\lambda_{\text{weak}} = 0.5$ with the highest DSC and the lowest MSD. On the clinical dataset $\mathcal{D}_c$, a smaller λ_{weak} caused the model to remain biased toward simulated defect statistics, whereas an excessively large weight such as 0.7 led to performance saturation or degradation due to over-emphasis on weak labels. Based on these findings, λ_{weak} was set to 0.5 for all subsequent experiments to ensure balanced adaptation between the simulation and clinical domains.

4 Conclusion

We presented Diff-SABR, a symmetry-aware latent diffusion framework for reconstructing alveolar bone defects from CBCT under limited clinical supervision. By combining hyperbolic-paraboloid defect simulation, weak clinical supervision,

Table 3. Sensitivity analysis of λ_{weak} on synthetic ($\mathcal{D}_s$) and clinical ($\mathcal{D}_c$) datasets.

λ_{weak}	$\mathcal{D}_s$					$\mathcal{D}_c$				
	DSC (↑)	Precision (↑)	Recall (↑)	HD95 (↓)	MSD (↓)	DSC (↑)	Precision (↑)	Recall (↑)	HD95 (↓)	MSD (↓)
0.1	0.89 ± 0.06	0.85 ± 0.10	$\mathbf{0.94 \pm 0.03}$	1.12 ± 1.17	0.19 ± 0.11	0.71 ± 0.10	0.72 ± 0.13	0.73 ± 0.11	2.38 ± 0.95	0.61 ± 0.25
0.3	0.89 ± 0.06	0.85 ± 0.09	0.93 ± 0.04	1.20 ± 1.10	0.20 ± 0.10	0.75 ± 0.05	0.75 ± 0.13	0.77 ± 0.11	2.33 ± 0.81	0.54 ± 0.15
0.5	$\mathbf{0.91 \pm 0.03}$	0.90 ± 0.05	0.91 ± 0.05	0.96 ± 0.68	0.16 ± 0.11	$\mathbf{0.82 \pm 0.04}$	$\mathbf{0.83 \pm 0.08}$	$\mathbf{0.82 \pm 0.06}$	$\mathbf{1.67 \pm 0.76}$	$\mathbf{0.36 \pm 0.08}$
0.7	0.91 ± 0.04	$\mathbf{0.92 \pm 0.04}$	0.89 ± 0.05	$\mathbf{0.85 \pm 0.87}$	$\mathbf{0.16 \pm 0.07}$	0.77 ± 0.03	0.77 ± 0.07	0.79 ± 0.07	1.79 ± 0.55	0.46 ± 0.09

and affine symmetry alignment via AlignNet, the proposed method improves anatomical fidelity and robustness on real clinical scans. Experimental results and ablations demonstrate that hybrid supervision and symmetry alignment are both essential to bridge the synthetic-to-clinical gap. Diff-SABR enables objective 3D reconstruction and can support preoperative graft volume estimation and implant treatment planning.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

Acknowledgments. This work was supported by the National Research Foundation of Korea (NRF) grants funded by the Korean government (MSIT) (RS-2023-NR077088 and RS-2026-25504135), and by the Technology Innovation Program (or Industrial Strategic Technology Development Program–Advanced Biomaterials) (RS-2025-14322975), funded by the Ministry of Trade, Industry & Energy (MOTIE).

References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: International conference on machine learning. pp. 214–223. Pmlr (2017)
2. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. Medical image analysis **12**(1), 26–41 (2008)
3. Chou, P.Y., Denadai, R., Hallac, R.R., Dumrongwongsiri, S., Hsieh, W.C., Pai, B.C., Lo, L.J.: Comparative volume analysis of alveolar defects by 3d simulation. Journal of clinical medicine **8**(9), 1401 (2019)
4. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: International conference on medical image computing and computer-assisted intervention. pp. 424–432. Springer (2016)
5. Cui, Z., Fang, Y., Mei, L., Zhang, B., Yu, B., Liu, J., Jiang, C., Sun, Y., Ma, L., Huang, J., et al.: A fully automatic ai system for tooth and alveolar bone segmentation from cone-beam ct images. Nature communications **13**(1), 2096 (2022)
6. Friedrich, P., Wolleb, J., Bieder, F., Thieringer, F.M., Cattin, P.C.: Point cloud diffusion models for automatic implant generation. In: International conference on medical image computing and computer-assisted intervention. pp. 112–122. Springer (2023)
7. Gu, S., Chen, D., Bao, J., Wen, F., Zhang, B., Chen, D., Yuan, L., Guo, B.: Vector quantized diffusion model for text-to-image synthesis. In: Proceedings of

the IEEE/CVF conference on computer vision and pattern recognition. pp. 10696–10706 (2022)
8. Guo, R., Ma, N., Shi, T., Wang, Y., Qin, Q., Chen, Z., Zuo, F., Wang, Y., Li, W.: Fully automatic ai diagnosis of alveolar bone defect in anterior teeth on cbct images. Clinical Oral Investigations **30**(1), 60 (2026)
9. Guo, Y., Pei, Y., Chen, S., Zhou, Z.b., Xu, T., Zha, H.: Stochastic anomaly simulation for maxilla completion from cone-beam computed tomography. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 667–677. Springer (2024)
10. Hu, C., Cao, N., Li, X., He, Y., Zhou, H.: Cbct-to-ct synthesis using a hybrid u-net diffusion model based on transformers and information bottleneck theory. Scientific Reports **15**(1), 10816 (2025)
11. Janssen, N.G., Schreurs, R., Bittermann, G.K., Borstlap, W.A., Koole, R., Meijer, G.J., Maal, T.J.: A novel semi-automatic segmentation protocol for volumetric assessment of alveolar cleft grafting procedures. Journal of Cranio-Maxillofacial Surgery **45**(5), 685–689 (2017)
12. Ji, X., Wang, W., Zhang, X., Chen, X.: Idnet: A diffusion model-enhanced framework for accurate cranio-maxillofacial bone defect repair. Bioengineering **12**(4), 407 (2025)
13. Jiang, Z., Chen, Y., Zhang, Y., Ge, Y., Yin, F.F., Ren, L.: Augmentation of cbct reconstructed from under-sampled projections using deep learning. IEEE transactions on medical imaging **38**(11), 2705–2715 (2019)
14. Linderup, B.W., Küseler, A., Jensen, J., Cattaneo, P.M.: A novel semiautomatic technique for volumetric assessment of the alveolar bone defect using cone beam computed tomography. The Cleft Palate-Craniofacial Journal **52**(3), 47–55 (2015)
15. Lu, K., Ren, L., Yin, F.F.: A geometry-guided deep learning technique for cbct reconstruction. Physics in Medicine & Biology **66**(15), 15LT01 (2021)
16. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
17. Su, J.Y., Chang, Y.C.: Complex alveolar bony defect reconstruction with three dimensional printing model to assist custom-made titanium mesh for guided bone regeneration. Journal of Dental Sciences **16**(2), 778 (2020)
18. Torabi, S., Panjnoush, M., Poursafar, F., Taheri, M., Khorsand, A., Yaghobee, S.: Comparison of ridge mapping and cone beam computed tomography for the determination of alveolar ridge width. Journal of Maxillofacial and Oral Surgery **21**(3), 802–807 (2022)
19. Vicente, A., Hung, K.F., Xu, Z., Yang, J., Li, J., Wiedel, A.P., Becker, M., Brogårdh-Roth, S., Ding, P., Hellén-Halme, K., et al.: Deep learning for automated alveolar cleft segmentation and bone graft volume estimation in cone-beam computed tomography imaging: a multicenter study. Oral surgery, oral medicine, oral pathology and oral radiology **141**(3), 400–408 (2026)
20. Wang, B., Liu, Z., Li, Y., Xiao, X., Zhang, R., Cao, Y., Cui, L., Zhang, P.: Cranial implant design using a deep learning method with anatomical regularization. In: Li, J., Egger, J. (eds.) Towards the Automatization of Cranial Implant Design in Cranioplasty. pp. 85–93. Springer International Publishing, Cham (2020)